

Claims

- [c1] In a computer network including at least two workstations and implementing a grid for executing a grid workload on the shared resources of the at least two workstations, a method of assigning said grid workload to one of the at least two workstations, the method comprising:
- determining an activity state of a first workstation of the at least two workstations, said activity state being indicative of a degree of current utilization of said first workstation by a local workload;
 - in response to a determination that said activity state indicates low utilization, assigning said grid workload to said first workstation;
 - in response to a determination that said activity state indicates high utilization, holding said grid workload for assignment to another of the at least two workstations.
- [c2] The method of claim 1 wherein, in response to a determination that said activity state indicates high utilization, the method further comprising:
- determining a workstation preference for accepting said grid workload on said first workstation when said utilization is high;
 - in response to a determination that said workstation preference is yes, assigning said grid workload to said first workstation;
 - in response to a determination that said workstation preference is no, holding said grid workload for assignment to another of the at least two workstations.
- [c3] The method of claim 2, further comprising:
- determining a delay level associated with said grid workload and indicative of an extent to which said grid workload will lock the resources of a workstation to which it is assigned, and;
- wherein, in response to a determination that said activity state indicates high utilization and that said workstation preference is yes, assigning said grid workload to said first workstation only if said delay level indicates that said grid workload will lock the workstation resources to an extent below a threshold.
- [c4] In a computer network including at least two workstations and implementing a

grid for executing a grid workload on the shared resources of the at least two workstations, a method of assigning said grid workload to one of the at least two workstations, the method comprising:

- monitoring a utilization level indicative of the extent to which a local workload is utilizing the resources of a first workstation of the at least two workstations and setting a first workstation activity state to;

- idle if said utilization level is below a first threshold, or to

- active if said utilization level is above said first threshold and said first workstation activity state is already idle or active, or to

- blocked if said utilization level is above said first threshold and said first workstation activity state is already blocked, and;

- checking said first workstation activity state;

- in response to a first workstation activity state of idle, assigning said grid workload to said first workstation;

- in response to a first workstation activity state of active, determining a first workstation preference for accepting said grid workload when said first workstation activity state is active;

- in response to a determination that said preference is yes, assigning said grid workload to said first workstation;

- in response to a determination that said preference is no;

- setting said first workstation activity state to blocked;

- holding said grid workload for assignment to another of the at least two workstations; and

- in response to a first workstation activity state of blocked, holding said grid workload for assignment to another of the at least two workstations.

[c5]

The method of claim 4, further comprising:

determining a delay level associated with said grid workload and indicative of the extent to which said grid workload will lock the resources of a workstation to which it is assigned, and;

wherein, in response to a first workstation activity state of active and a first workstation preference of yes, assigning said grid workload to said first

workstation only if said delay level indicates that the grid workload will lock the

workstation resources to an extent below a second threshold.

[c6] In a computer network including at least two workstations and implementing a grid for executing a grid workload on the shared resources of the at least two workstations, a method of assigning said grid workload to one of the at least two workstations, the method comprising:

- running a hypervisor on each of the at least two workstations;
- executing a local workload on a first virtual machine running under each of said hypervisors, each hypervisor giving priority to said each first virtual machine;
- running at least one additional virtual machine under each hypervisor, said each at least one additional virtual machine executing an operating system capable of operating as part of said grid;
- monitoring a utilization of said each workstation by said each local workload and determining a local workload activity level indicative thereof;
- in response to a determination that said local workload activity level indicates said utilization is below a first threshold, sharing the resources of said workstation with said at least one additional virtual machine in order to execute said grid workload.

[c7] The method of claim 6 wherein said determination of said each local workload activity level further comprises classifying said each local workload activity level into one of at least three categories, the at least three categories including the categories of idle, active and blocked, wherein

- said idle category indicates said utilization is below said first threshold,
- said active category indicates said utilization is above said first threshold and a workstation preference is that said workstation resources may be shared with said at least one additional virtual machine when said utilization is above said first threshold, and
- said blocked category indicates said utilization is above said first threshold and said workstation preference is that said workstation resources may not be shared with said at least one additional virtual machine when said utilization is above said first threshold, and;

wherein, said workstation resources are only shared with said at least one additional virtual machine in order to execute said grid workload if said local

workload activity level is idle or active.

[c8] The method of claim 7 further comprising:
determining a delay level associated with said grid workload and indicative of the extent to which said grid workload will lock the resources of a workstation to which it is assigned, and;
wherein, in response to said local workload activity level being active and said workstation preference indicating that said workstation resources may be shared with said at least one additional virtual machine when said utilization is above said first threshold, sharing said workstation resources with said at least one additional virtual machine only if said delay level indicates that said grid workload will lock said workstation resources to an extent below a second threshold.

[c9] In a computer network including at least two workstations and implementing a grid for executing a grid workload on the shared resources of the at least two workstations, a system for assigning said grid workload to one of the at least two workstations, the system comprising:
– a workstation monitor for determining an activity state of a first workstation of the at least two workstations, said activity state being indicative of a degree of current utilization of said first workstation by a local workload; and
– a workload scheduler for;
in response to a determination that said activity state indicates low utilization, assigning said grid workload to said first workstation; or
in response to a determination that said activity state indicates high utilization, holding said grid workload for assignment to another of the at least two workstations.

[c10] The system of claim 9 wherein,
– said workstation monitor further determines a workstation preference for accepting said grid workload on said first workstation when said utilization is high; and
– said workload scheduler, in response to a determination that said activity state indicates high utilization, further;

- checks said workstation preference, and, in response to a workstation preference of yes, assigns said grid workload to said first workstation, or, in response to a workstation preference of no, holds said grid workload for assignment to another of the at least two workstations.

[c11] The system of claim 10 further comprising:

- a grid workload monitor for determining a delay level associated with said grid workload and indicative of the extent to which said grid workload will lock the resources of a workstation to which it is assigned, and wherein;
- said workload scheduler, in response to a determination that said activity state indicates high utilization and that said workstation preference is yes, further assigns said grid workload to said first workstation only if said delay level indicates that said grid workload will lock the workstation resources to an extent below a threshold.

[c12] In a computer network including at least two workstations and implementing a grid for executing a grid workload on the shared resources of the at least two workstations, a system for assigning said grid workload to one of the at least two workstations, the system comprising:

- a workstation monitor for:
 - monitoring a utilization level indicative of the extent to which a local workload is utilizing the resources of a first workstation of the at least two workstations and setting a first workstation activity state to;
 - idle if said utilization level is below a first threshold, or to
 - active if said utilization level is above said first threshold and said first workstation activity state is already idle or active, or to
 - blocked if said utilization level is above said first threshold and said first workstation activity state is already blocked, and;
 - determining a first workstation preference for accepting said grid workload when said first workstation activity state is active; and
- a workload scheduler for:
 - checking said first workstation activity state;
 - in response to a workstation activity state of idle, assigning said grid workload to said first workstation;

- in response to a first workstation activity state of active, checking said first workstation preference;
 - in response to a first workstation preference of yes, assigning said grid workload to said first workstation;
- in response to a first workstation preference of no;
 - setting said first workstation activity state to blocked;
 - holding said grid workload for assignment to another of the at least two workstations; and
- in response to a first workstation activity state of blocked, holding said grid workload for assignment to another of the at least two workstations.

[c13] The system of claim 12, further comprising:

- a grid workload monitor for determining a delay level associated with said grid workload and indicative of the extent to which said grid workload will lock the resources of a workstation to which it is assigned, and wherein; said workload scheduler, in response to a first workstation activity state of active and a first workstation preference of yes, assigns said grid workload to said first workstation only if said delay level indicates that said grid workload will lock the workstation resources to an extent below a second threshold.

[c14] In a computer network including at least two workstations and implementing a grid for executing a grid workload on the shared resources of the at least two workstations, a system for assigning said grid workload to one of the at least two workstations, the system comprising:

- a hypervisor running on each of the at least two workstations;
- a local workload executing on a first virtual machine running under each of said hypervisors, said each hypervisor giving priority to said each first virtual machine;
- at least one additional virtual machine running under said each hypervisor, said each at least one additional virtual machine executing an operating system capable of operating as part of said grid;
- a workstation monitor for monitoring a utilization of said each workstation by said each local workload and determining a local workload activity level indicative thereof;

– a workload scheduler which, in response to a determination that said local workload activity level indicates said utilization is below a first threshold, shares the resources of said workstation with said at least one additional virtual machine in order to execute said grid workload.

[c15] 1)The system of claim 14 wherein said workstation monitor further classifies said each local workload activity level into one of at least three categories, the at least three categories including the categories of idle, active and blocked, wherein

- said idle category indicates said utilization is below said first threshold,
- said active category indicates said utilization is above said first threshold and a workstation preference is that said workstation resources may be shared with said at least one additional virtual machine when said utilization is above said first threshold, and
- said blocked category indicates said utilization is above said first threshold and said workstation preference is that said workstation resources may not be shared with said at least one additional virtual machine when said utilization is above said first threshold, and wherein;

said workload scheduler shares said workstation resources with said at least one additional virtual machine in order to execute said grid workload only if said local workload activity level is idle or active.

[c16] The system of claim 15 further comprising:

- a grid workload monitor for determining a delay level associated with said grid workload and indicative of the extent to which said grid workload will lock the resources of a workstation to which it is assigned, and wherein;

said workload scheduler, in response to said local workload activity level being active, shares said workstation resources with said at least one additional virtual machine only if said delay level indicates that said grid workload will lock said workstation resources to an extent below a second threshold.